Parallel multi-grid summation for the N-body problem

Jesús A. Izaguirre
with Thierry Matthey

Department of Computer Science and Engineering
University of Notre Dame

October 2, 2003

This work is partially supported by two NSF grants (CAREER and BIOCOMPLEXITY) and two grants from University of Notre Dame

Talk Roadmap

1. The N-body problem

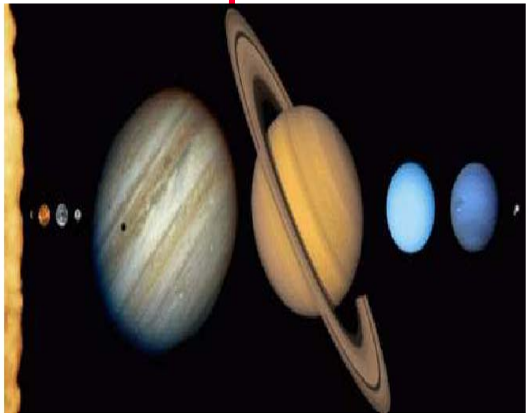

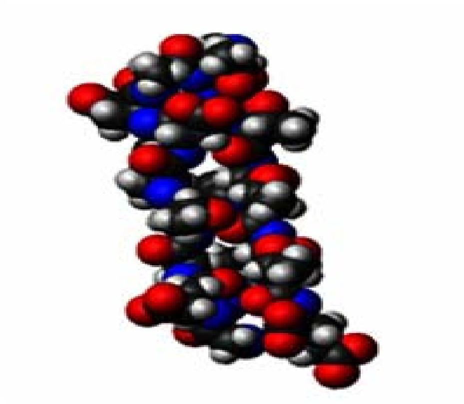

2. Multilevel methods for N-body problem

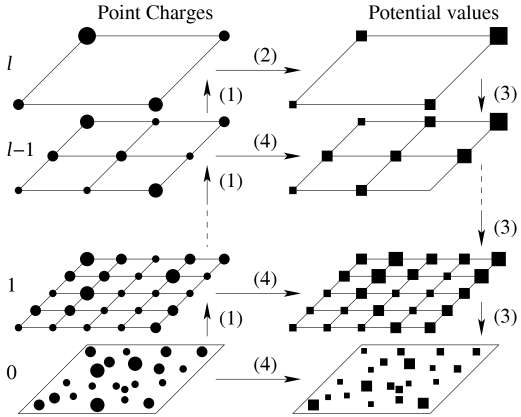

3. Evaluation of parallel multilevel methods

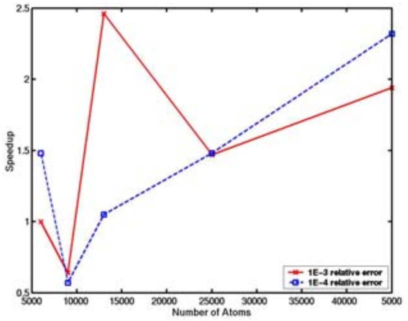

4. Discussion and extensions

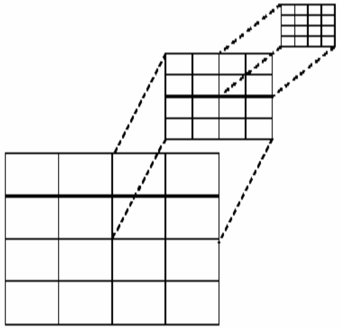

The N-body problem I

The motion of planets in the solar system is an N-body problem, described by Newton in 1687 in the *Philosophiae Naturalis Principia Mathematica.*

$$\vec{F}_{12} = -G\frac{m_1 m_2}{||\vec{r}_1 - \vec{r}_2||^2}\frac{\vec{r}_1 - \vec{r}_2}{||\vec{r}_1 - \vec{r}_2||}$$

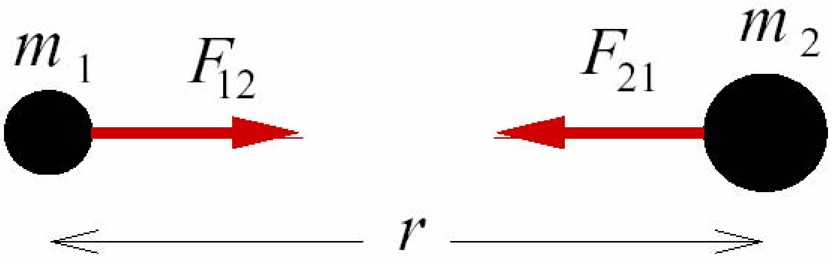

The N-body problem II

- Laplace (1820) dreamed about predicting long term motion of bodies in nature
- For three or more bodies, Poincaré (1900s) discovered that Newtonian mechanics does not have an analytical solution, and that it is chaotic
- For relatively small systems and short time trajectories, it can be solved numerically
- For larger systems or long time trajectories, one can gather important *statistics*

The N-body problem III

In the N-body problem one computes the interaction forces for particles or bodies, and then computes the motion due to that force using the celebrated Newton's second law:

$$m_i \frac{\mathrm{d}^2}{\mathrm{d}t^2}\vec{r}_i(t) = \vec{F}_i(t),$$

where the mass of the i^{th} atom is m_i. Its instantaneous force $\vec{F}_i(t)$ is typically a conservative force,

$$\vec{F}_i(t) = -\nabla_i U(\vec{r}_1(t), \vec{r}_2(t), \ldots, \vec{r}_N(t)), \ \ i = \{1, 2, \ldots, N\}.$$

The N-body problem IV

- Examples of N-body problems include:
  - the motion of planets and galaxies
  - the folding of proteins
  - electronic structure of materials

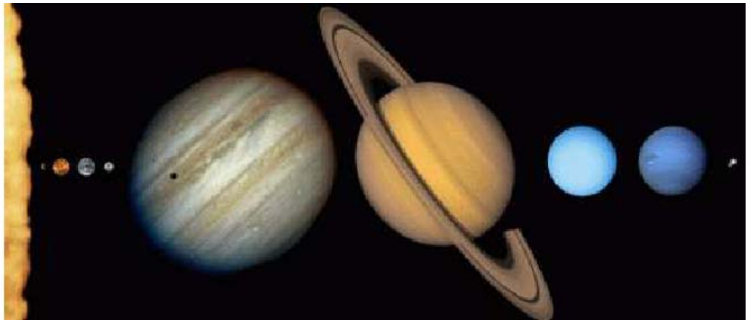

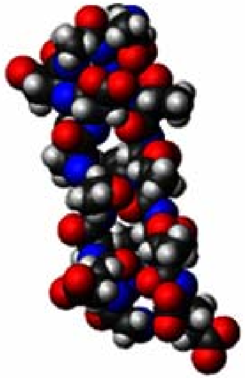

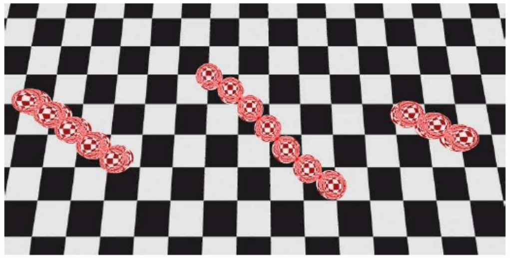

Ugelstad spheres in a magnetic fluid and an applied magnetic field.

The N-body problem V

More formally, we consider the problem of electrostatic interactions for isolated and periodic systems

Vacuum

The N-body problem for an isolated system with N particles can be expressed by

$$U^{\text{electrostatic}}(\vec{r}_1, \vec{r}_2, \ldots, \vec{r}_N) = \frac{1}{2} \sum_{i=1}^{N} \sum_{j \notin \chi(i)} \frac{q_i q_j}{4\pi\varepsilon_0 |\vec{r}_j - \vec{r}_i|},$$

where the position and partial charge of the i^{th} atom are $\vec{r}_i$ and q_i, and the dielectric coefficient is ε_0.

The N-body problem VI

Periodic boundary conditions

In the case of periodic boundary conditions, the problem to be solved is

$$U^{\text{electrostatic}}(\vec{r}_1, \vec{r}_2, \ldots, \vec{r}_N) = \frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{\vec{m}' \in \mathbb{Z}^3} \frac{q_i q_j}{4\pi\varepsilon_0 |\vec{r}_j - \vec{r}_i + \vec{m}L|},$$

where the sum is over all periodic cells with index $\vec{m}$, with self-interactions excluded, and L is the length of a periodic box cell with dimensions $L \times L \times L$. The primed sum also excludes interactions for pairs that are in the exclusion list. This is a conditionally convergent sum.

Talk Roadmap

1. The N-body problem

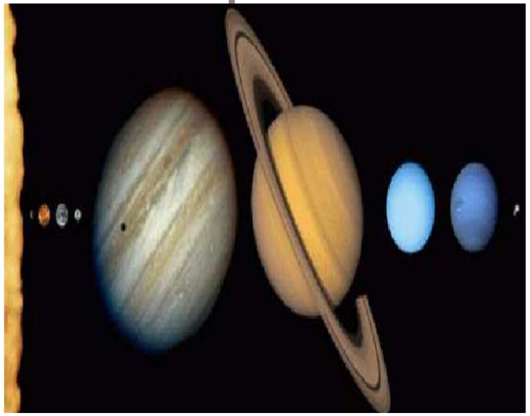

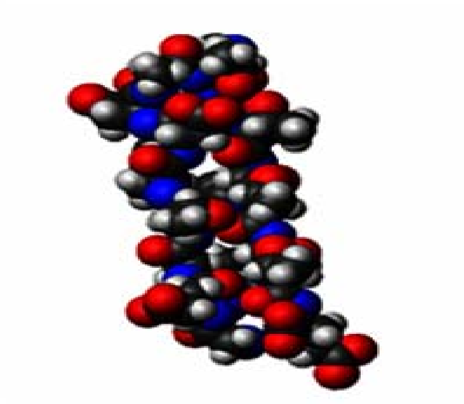

2. Multilevel methods for N-body problem

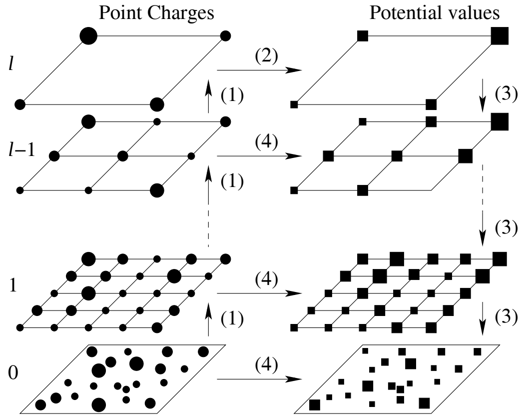

3. Evaluation of parallel multilevel methods

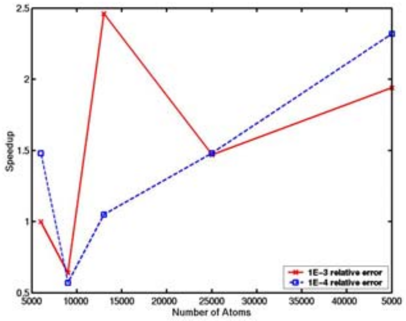

4. Discussion and extensions

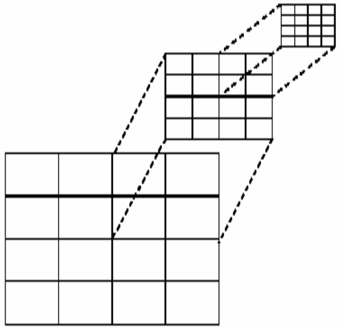

Outline

- Describe multilevel methods that solve the N-body problem efficiently, in *O(N)* operations, both for isolated and periodic systems
- Compare to other solution methods
- Show scalable parallel implementation of this multilevel method
- Show how to determine optimal parameters for this algorithm
- Discuss limitations, extensions, and future work

Comparison of fast electrostatic algorithms

Direct, particle-particle	$\Theta(N^2)$	*Laplace, 1820*
Ewald	$\Theta(N^{3/2})$	*Ewald, 1921*
Barnes-Hut	$\Theta(N \log N)$	*Barnes & Hut, 1986*
Particle Mesh Ewald	$\Theta(N \log N)$	*Darden, 1993*
Fast multipole	$\Theta(N)$	*Greengard & Rokhlin, 1987*
Multigrid summation	$\Theta(N)$	*Brandt et al., 1990* *Skeel et al., 2002* *Izaguirre & Matthey, 2003*

Isolated Systems I: Particle-Particle (PP) or direct method

Particle-Particle (PP):

- Exact evaluation for isolated systems
- Prohibitive cost: $\Theta(n^2)$

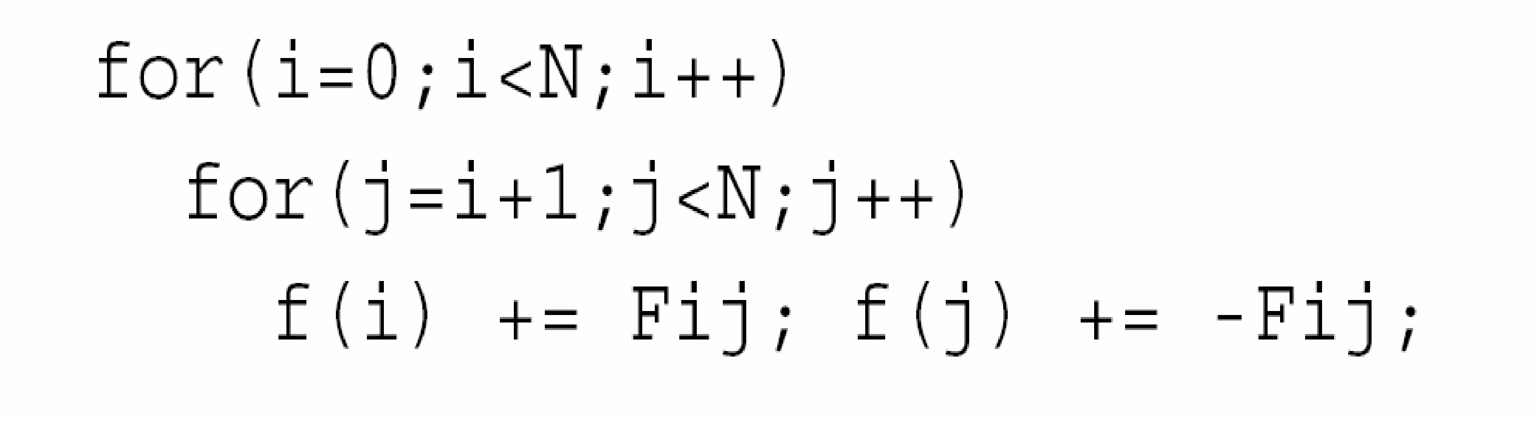

```
for(i=0;i<N;i++)
  for(j=i+1;j<N;j++)
    f(i) += Fij; f(j) += -Fij;
```

Periodic Systems: Ewald summation I

- To solve conditionally convergent sum, Ewald discovered how to split the sum into two sums
- This idea of splitting is the basis of fast electrostatics methods

$$\frac{1}{r} = \underbrace{\left(\frac{1}{r} - g^{1}_{\text{smooth}}(r)\right)}_{\text{short-ranged}} + \underbrace{g^{1}_{\text{smooth}}(r)}_{\text{smooth}}$$

Ewald Summation II

- Atoms pairs in the direct lattice
  - Direct calculations
- Atoms pairs in the images
  - Fourier transforms

    A form of trigonometric interpolation, used when dealing with cyclic phenomena.

Ewald Summation III: Ewald's trick

Two concepts used here:

- *Split the problem:* Split the summation of electrostatic potential energy into two absolutely and rapidly converging series in direct and reciprocal space
- *Approximate a coarser subproblem:* trigonometric interpolation of the reciprocal space

Reciprocal space - the orthogonal system associated with the atoms in the unit cell

Ewald Summation IV

Ewald chooses

$$g(r) = \operatorname{erfc}(\alpha r) = \frac{2}{\sqrt{\pi}} \int_{\alpha}^{\infty} \exp(-s^2)\,ds.$$

The short-range part is again solved directly, whereas the smooth part is solved by a Fourier series, which approximately looks like this:

$$\frac{1}{2} \sum_{\vec{k} \neq 0} \frac{1}{\vec{k}^2} \exp\left(\frac{-\vec{k}^2}{4\alpha^2}\right) \left| \underbrace{\sum_j q_j \exp\left(-i\vec{k} \cdot \vec{r}_j\right)}_{\hat{\rho}(\vec{k})} \right|.$$

Periodic Systems: Particle-Mesh

Particle-Mesh (PM):

- Evaluation of forces via mesh and FFT
- FFT has high communication requirements when parallelized
- Moderate cost: $\Theta(n \lg n)$

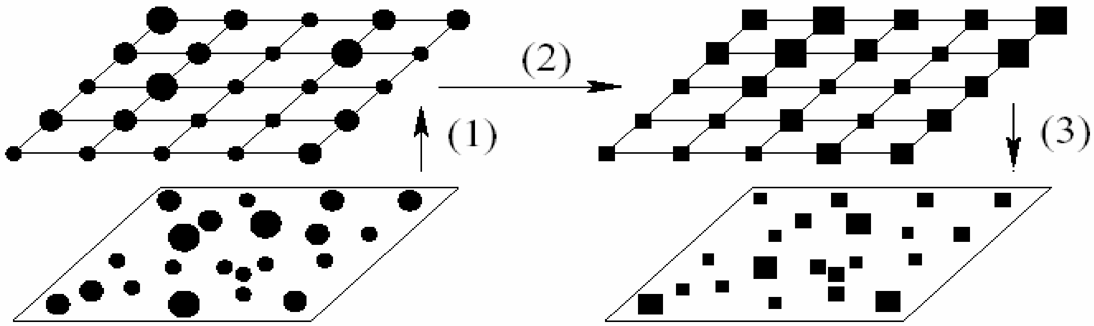

Periodic Systems: Particle Mesh Ewald

PME:

- α controls computation in PP or PM
- It can be chosen optimally so that

the complexity of the Ewald sum is $\Theta(N^{3/2})$ [cite: Finc94,BaCS01].

- The particle mesh Ewald (PME) method [cite: DARD93] chooses α so PP part is $\Theta(N)$, and uses $\Theta(N \log N)$ FFT for the PM part:

$$\hat{\rho}(\vec{k}) \approx \sum_n \exp(-i\vec{k} \bullet \vec{r}_n^h) \sum_j \phi_n(F_j) q_j,$$

where $\phi_n(F_j)$ project particle charges to the grid, and $\vec{r}_n^h$ projects particle positions to the grid.

Periodic or Isolated Systems: Multigrid Summation I

Multi-grid summation

1. Splits the computational kernel into short-range and smooth parts at a cutoff r_c using switching functions
2. Coarsening is done by interpolation from particles into a grid (PM) and then within a hierarchy of grids (MM)
3. $\Theta(N)$
4. Parallelized in program ProtoMol

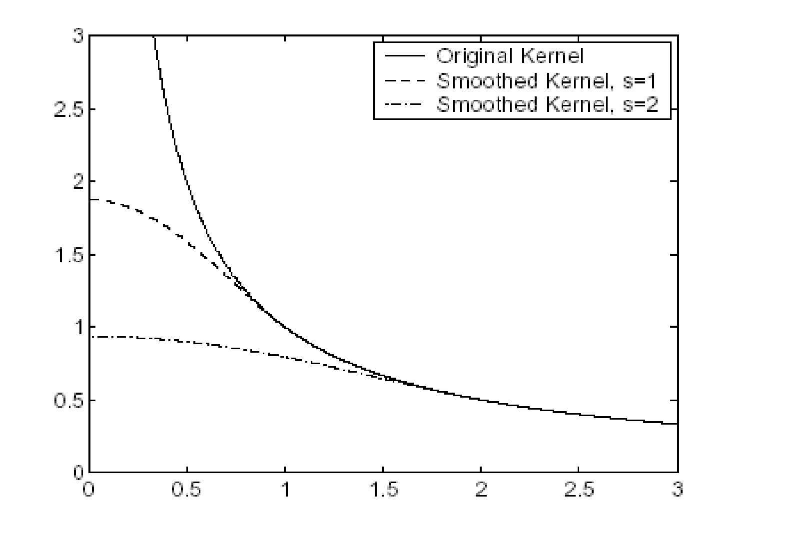

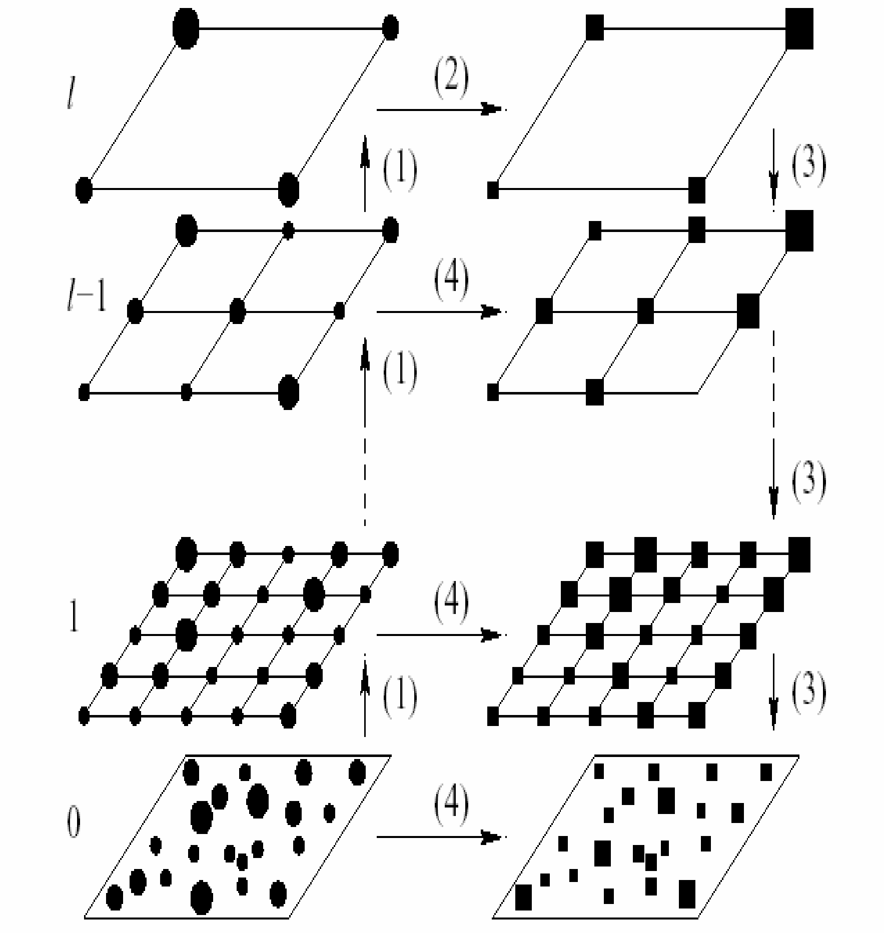

Multigrid Summation II

Separation of length scales

The separation of the length scales is done with a switching function that brings the computational kernel smoothly to zero at a cutoff distance r_{c}. These switching functions have varying degrees of smoothness and computational cost.

Coarsening

The softened kernel g^1_{smooth} is approximated at the source $\vec{r}'$:

$$g^1_{\mathrm{smooth}}(\|\vec{r}-\vec{r}'\|) \approx \sum_k g^1_{\mathrm{smooth}}(\|\vec{r}-\vec{r}_{h,k}\|)\phi_k(\vec{r}'),$$

where $\vec{r}_{h,k}$ are points on a 3-d grid with grid point separation h, and ϕ_k are piecewise polynomials with local support on a few grid cells.

Multigrid Summation III

The coefficients of the basis functions $g^1_{\text{smooth}}(\|\vec{r}-\vec{r}_{h,k}\|)$ are approximated at the destination $\vec{r}$:

$$g^1_{\text{smooth}}(\|\vec{r}-\vec{r}_{h,k}\|) \approx \sum_m g^1_{\text{smooth}}(\|\vec{r}_{h,m}-\vec{r}_{h,k}\|)\phi_m(\vec{r}),$$

resulting in a double sum over the grid cells necessary for the interpolation using ϕ_m and ϕ_k.

$$g^1_{\text{smooth}}(\|\vec{r}-\vec{r}'\|) \approx \sum_k \sum_m \phi_m(\vec{r}) g^1_{\text{smooth}}(\|\vec{r}_{h,m}-\vec{r}_{h,k}\|)\phi_k(\vec{r}').$$

Multigrid Summation IV

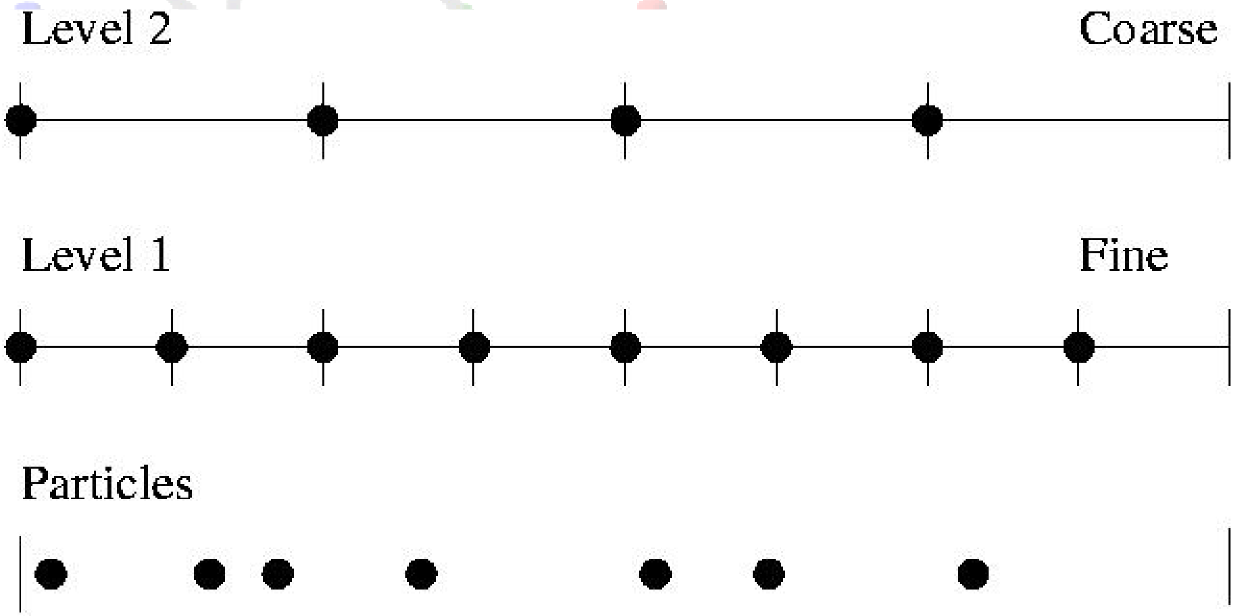

Multigrid Summation V

The charges at the grid point are defined as

$$q_{h,k} = \sum_{i=1}^{N} q_i \phi_k(\vec{r}_i).$$

Using these definitions, the smooth part of the electrostatic energy can be written simply as

$$\sum_{k} \sum_{m} q_{h,m} q_{h,k} \| \vec{r}_{h,m} - \vec{r}_{h,k} \|.$$

The sum over particle pairs has been reduced to a sum over grid point pairs.

Multigrid Summation VI

Hierarchical decomposition

The *smoothed kernel* $g^1_{\text{smooth}}(r)$ is further smoothed on a coarser grid using $g^2_{\text{smooth}}(r)$. The superscript indicates the grid level. Spliting $g^1_{\text{smooth}}(r)$:

$$g^1_{\text{smooth}}(r) = (g^1_{\text{smooth}}(r) - g^2_{\text{smooth}}(r)) + g^2_{\text{smooth}}(r),$$

where $g^1_{\text{smooth}}(r) - g^2_{\text{smooth}}(r)$ is zero for $r > 2r_{\text{c}}$. This process can be applied recursively for grid level $k \in \{1, 2, \ldots, l\}$.

Talk Roadmap

1. The N-body problem

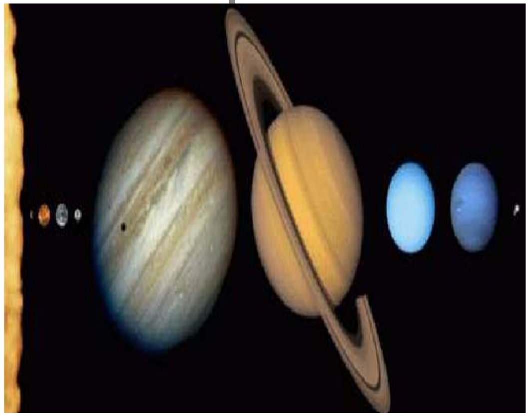

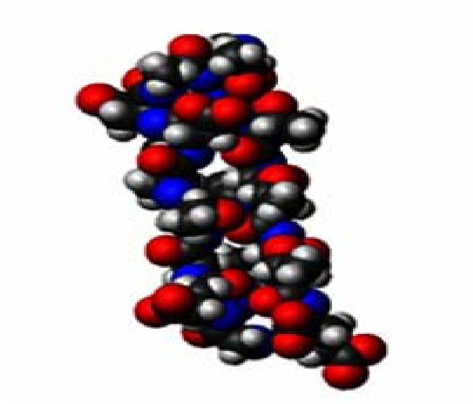

2. Multilevel methods for N-body problem

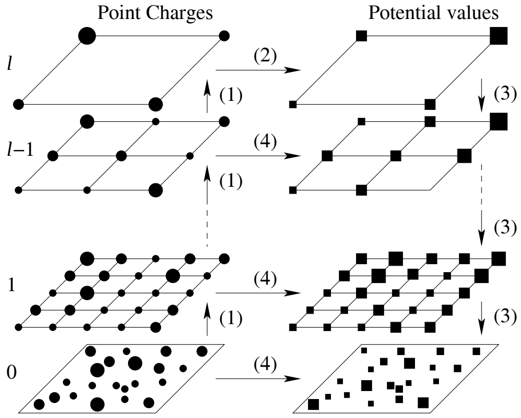

3. Evaluation of parallel multilevel methods

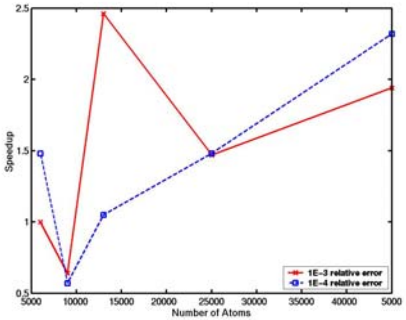

4. Discussion and extensions

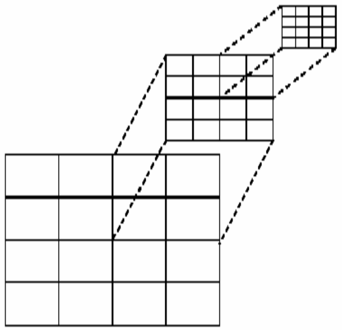

Parallelization of Multigrid summation I

- The data is distributed using a *d*-dimensional decomposition
- The same point across all grids belongs to the same processor
- Direct and smooth parts are parallelized at each grid

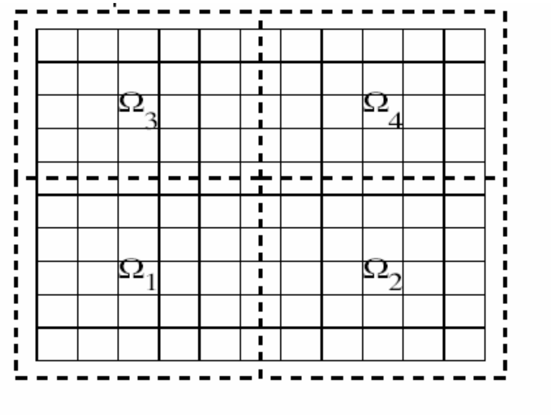

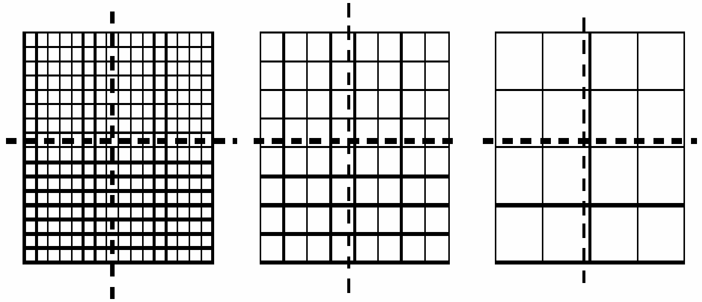

Grid h

Grid 2h

Grid 4h

Parallelization of Multigrid summation II

- Work distribution uses either a dynamic master-slave scheduling, or a static distribution

- Uses MPI and global communication

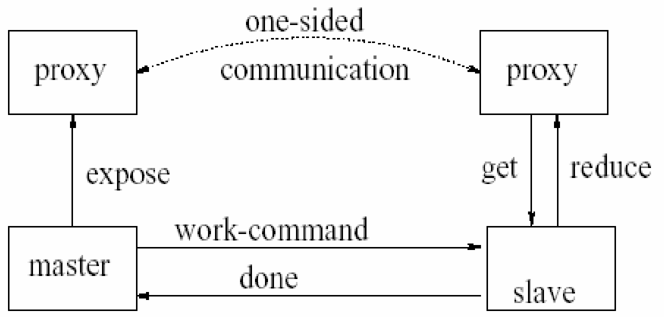

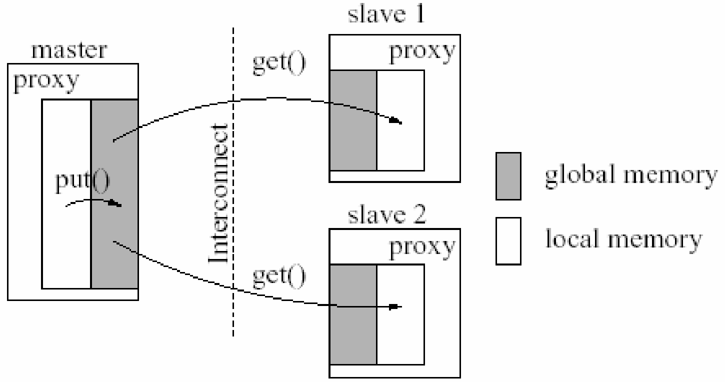

Experimental protocol

- These methods were tested and implemented in a common generic and OO framework, ProtoMol:
  1. Smooth Particle Mesh Ewald
  2. Multigrid summation
  3. Ewald summation
- Testing protocol:
  - Methods (1) and (2) above were compared against (3) to determine accuracy and relative speedup
  - Tested on atomic systems ranging from 1,000 to 1,000,000 atoms, and low and high accuracies
  - Optimal parameters for (1)-(3) determined using performance model and run-time testing
  - Tested on shared memory computer (SGI Altix) and Beowulf cluster (IBM Regatta)

Software adaptation I

(1) Domain specific language to define MD simulations in ProtoMol (Matthey and Izaguirre, 2001-3)

(2) "JIT" generation of prototypes: factories of template instantiations

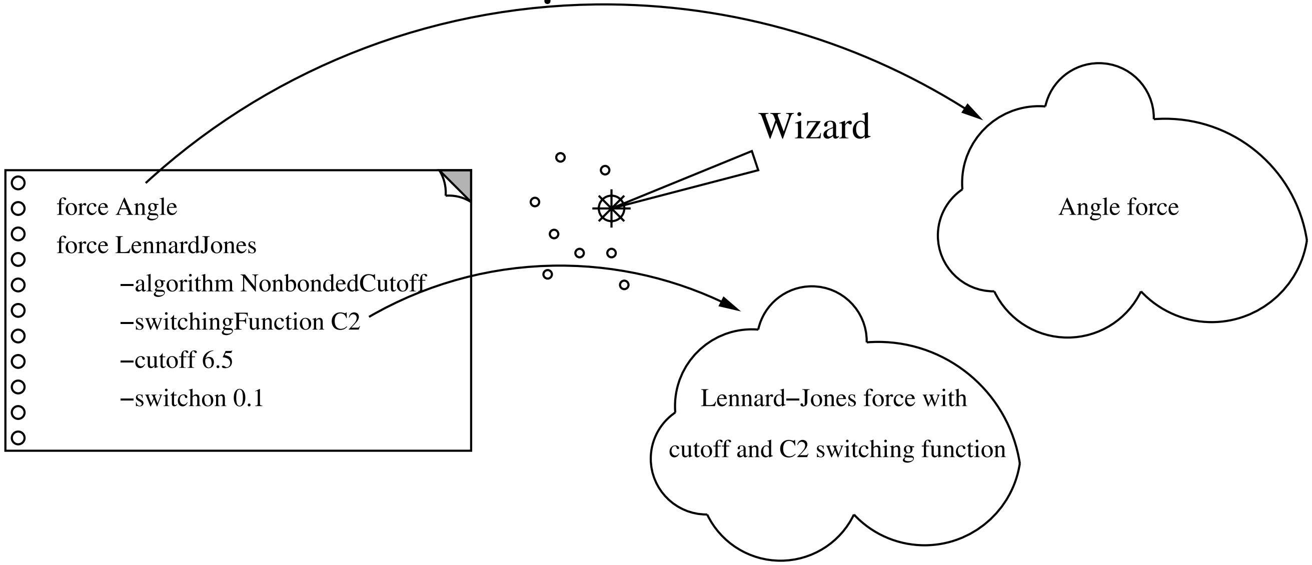

Software adaptation II

(3) Timing and comparison facilities:

```
force compare time force Coulomb -algorithm PMEwald
   -interpolation BSpline  -cutoff 6.5 -gridsize 10 10 10
force compare time force Coulomb -algorithm PMEwald
   -interpolation BSpline  -cutoff 6.5 -gridsize 20 20 20
```

(4) Selection at runtime – extensible module using Python

```
trial_param1[i]=cutoffvalue+1
test_PME(cutoffvalue+1,gridsize,...,
             accuracy)
```

Optimization strategies I

(1) Rules generated from experimental data & analytical models (thousands of data points)

(2A) At run-time, *M* tests are tried by exploring the parameter that changes accuracy or time most rapidly (*biased search),* or

(2B) At run-time, *M* tests are tried by randomly exploring all *valid* method parameter combinations

(3) Choose the fastest algorithm/parameter combination within accuracy constraints

Sequential performance of N-body solvers

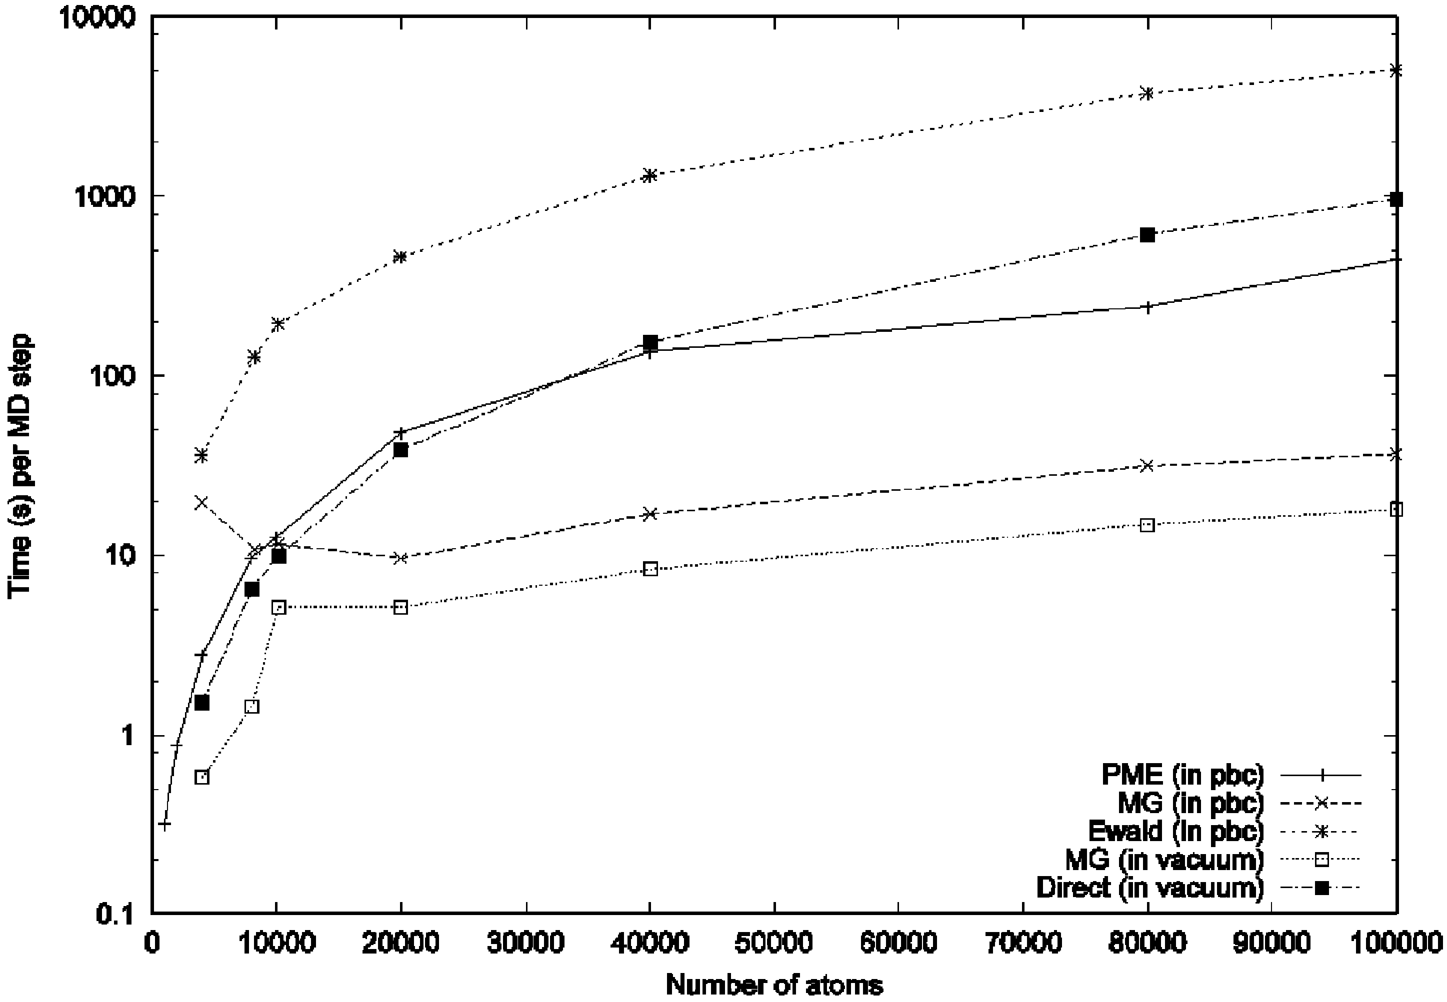

Fastest algorithms found

Platform: Linux; Biased search with 3 trials and unbiased with 5 trials

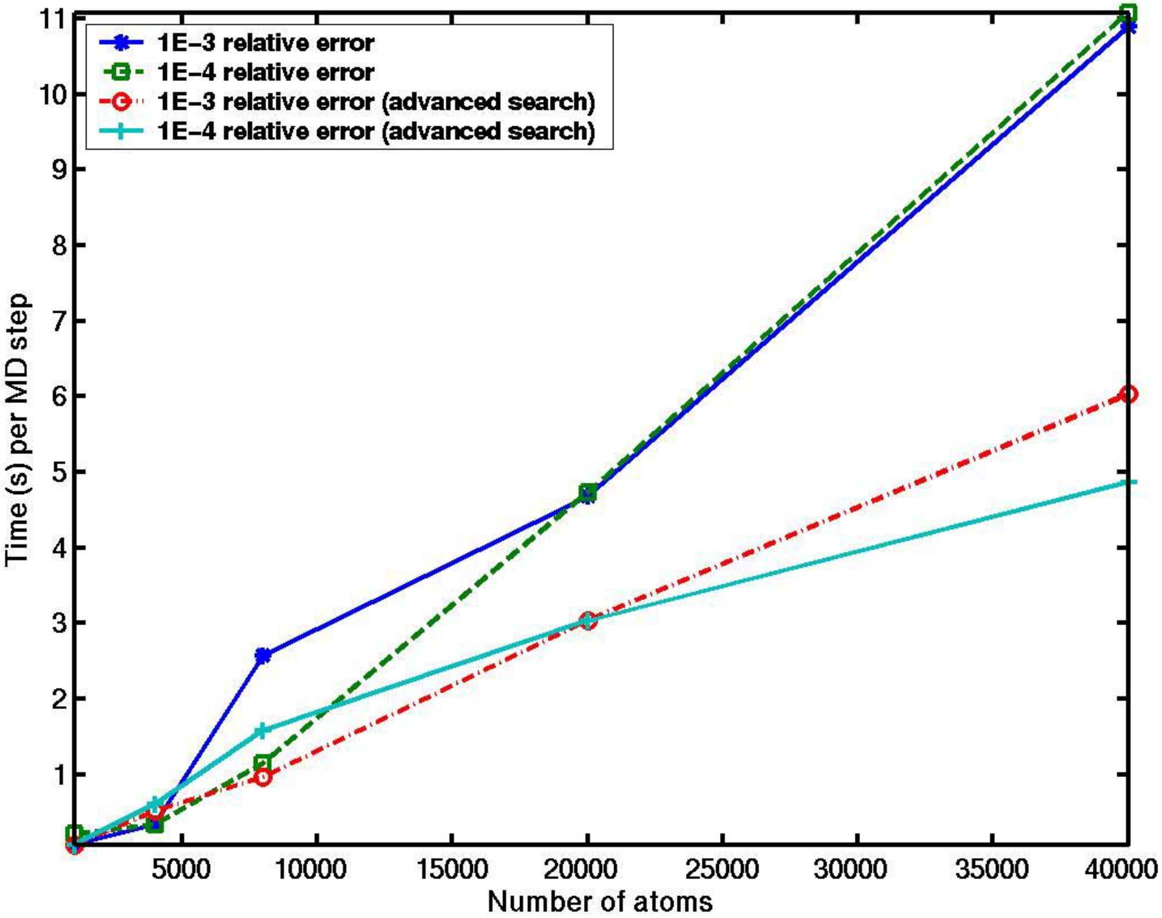

Parallel scalability of Multigrid summation

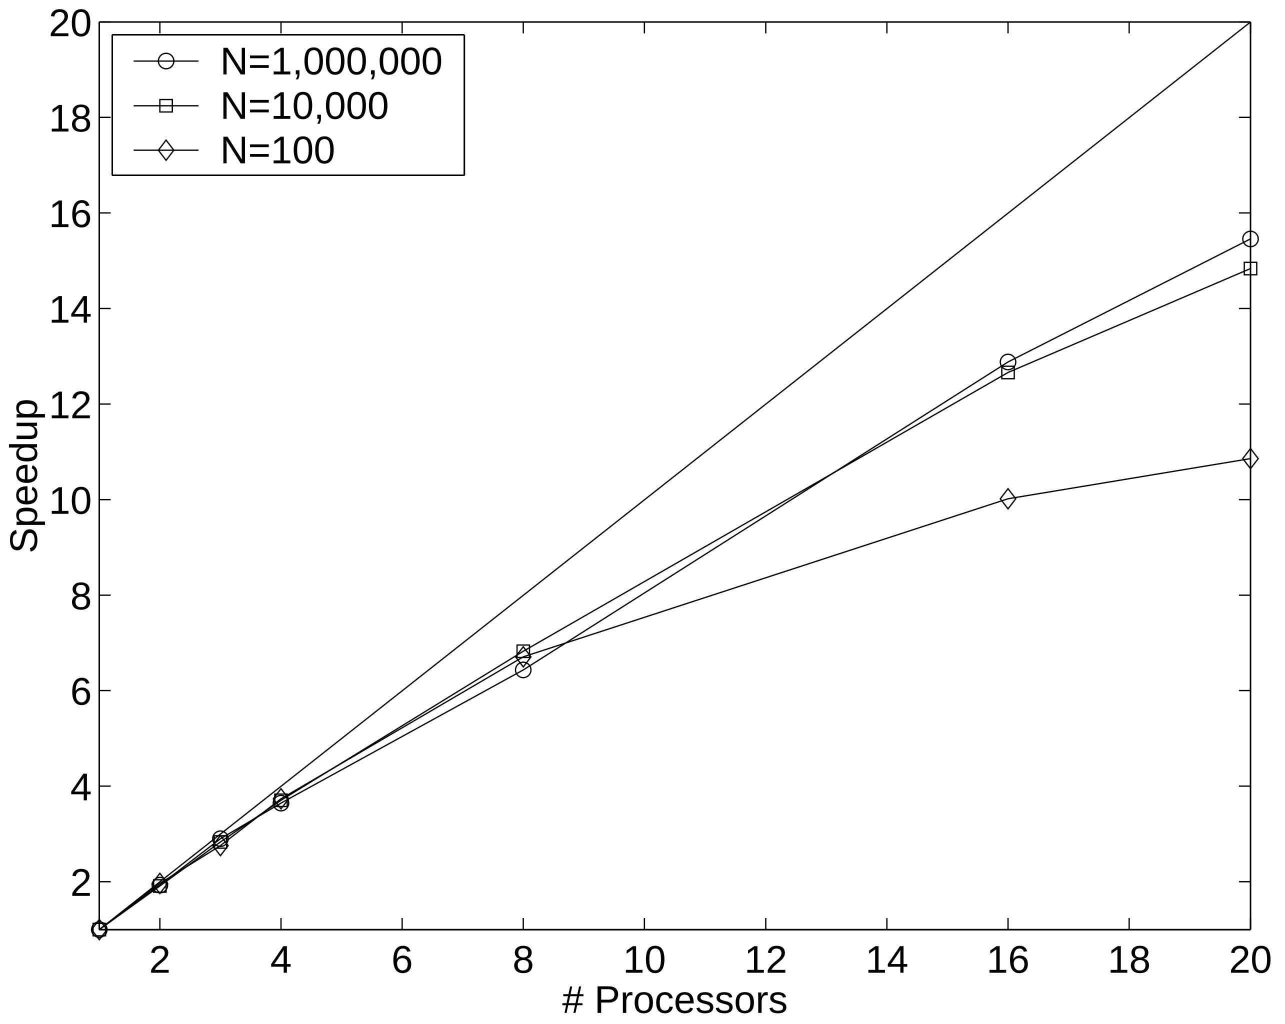

Parallel efficiency of Multigrid summation

Platform	**static**	**dynamic**
Altix ($p = 8$)	83%	84%
Regatta ($p = 10$)	87%	82%
Regatta ($p = 20$)	NA	75%

Parallel efficiency of Multi-grid on a shared memory Itanium 2 SGI Altix and a Linux Beowulf IBM Regatta cluster for 1,000,000 particles

Talk Roadmap

1. The N-body problem

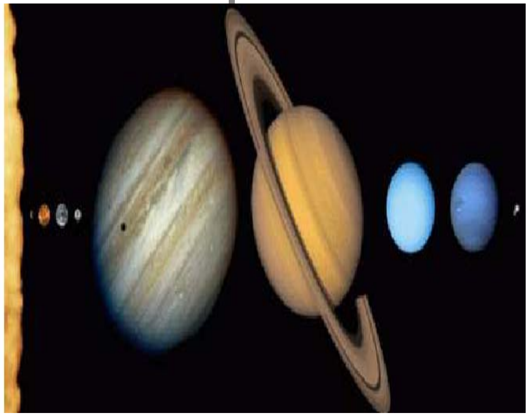

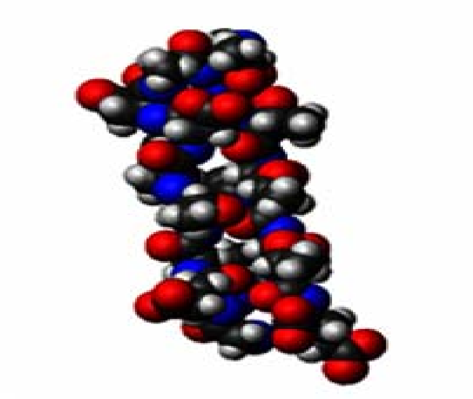

2. Multilevel methods for N-body problem

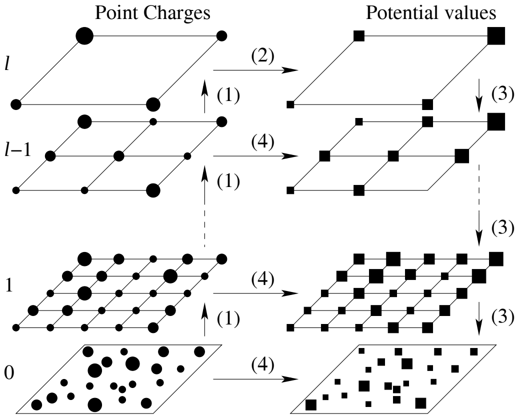

3. Evaluation of parallel multilevel methods

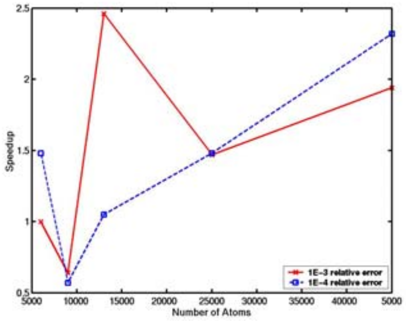

4. Discussion and extensions

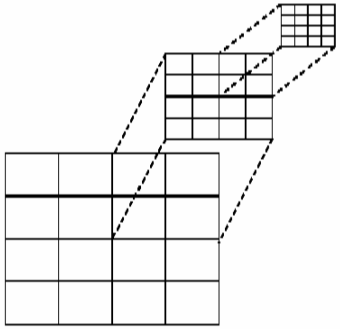

Summary

- Multigrid summation has been extended to periodic boundary conditions. It has been optimized and parallelized in open source program ProtoMol
- The key idea is to do a hierarchical decomposition of the computational kernel. Interpolation is used to approximate the kernel at coarser levels.
- It enables faster simulations of larger systems compared to other methods (1,000,000 particles), for example, Coulomb Crystals by Matthey *et al.*, 2003
- Optimization of parameters through web service MDSimAid: Ko & I., 2002; Wang, Crocker, Nyerges, & I. 2003)

Other methods: Barnes-Hut or Tree method

Barnes-Hut

- Cells are considered to be separated if the diameter of the cell is much smaller than the distance to a particle
- Coarsens at the source only
- $\Theta(N \log N)$
- Useful in astrophysics, due to non-uniformity of problem

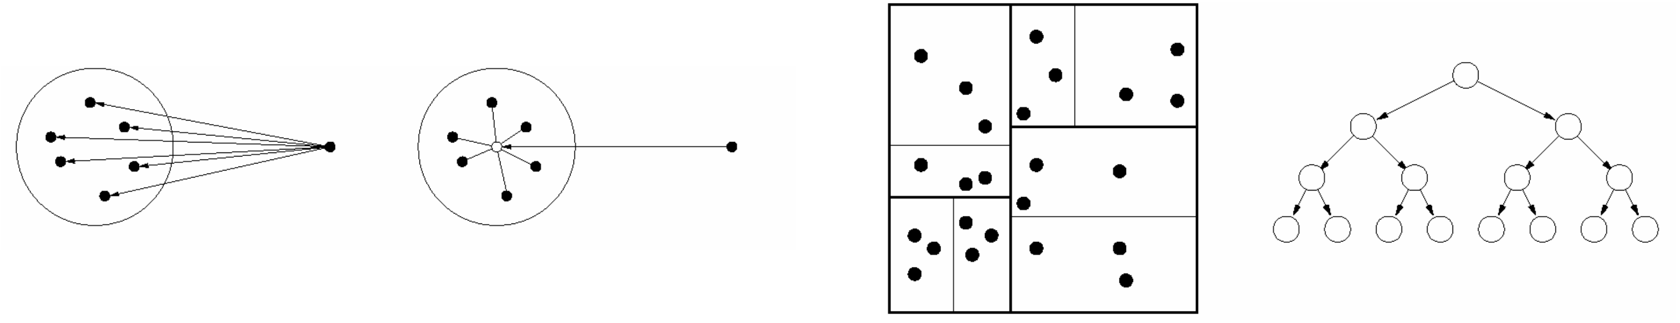

Other methods: Cell methods or Fast Multipole

Fast Multipole

- The effect of every charged particle is split into long and short range
- At short range, it is solved using PP
- At long range, it is represented by a Taylor expansion of spherical harmonic polynomials
- $\Theta(N)$
- Parallelized in program DPMTA

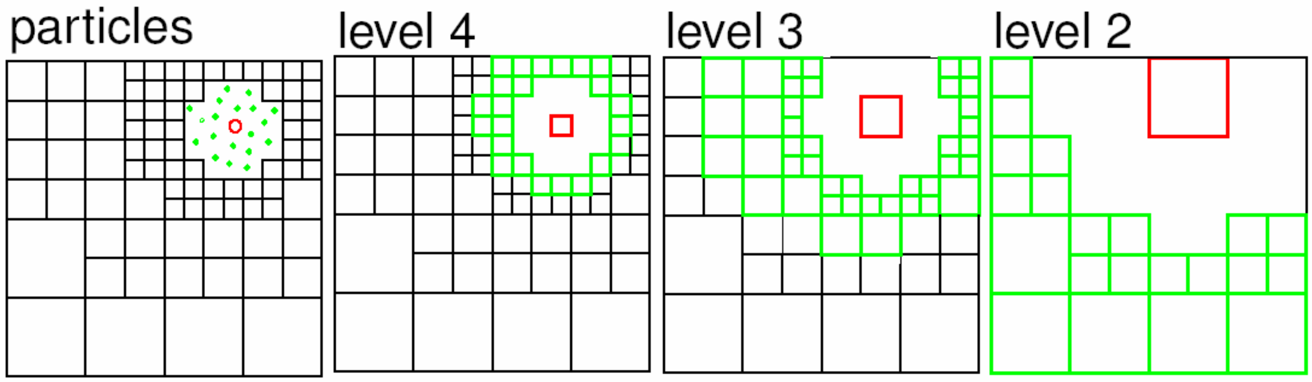

Discussion I

- Main problem in parallelization is that there is less work for all processors as the grid becomes coarser:
    - Agglomerate coarser grids
    - One could parallelize across grids as well - needs more sophisticated scheduling
    - Adaptive refinement would be useful for non-uniform systems (e.g., galaxies, sub-cellular systems)

Discussion II

- Low level parallelism can be optimized at run time as well
  - e.g., one-sided vs. two-sided communication; broadcast vs. point-to-point global communication
- Method efficient for low accuracy simulations. Need to improve it for higher accuracy
  - cleverer interpolation schemes

Acknowledgments

- Graduate students
  - Alice Ko
  - Yao Wang
- REU
  - Michael Crocker
  - Matthew Nyerges
- For further reference:
  - http://mdsimaid.cse.nd.edu
  - http://protomol.sourceforge.net
  - http://www.nd.edu/~izaguirr